

METHOD, SYSTEM, AND PROGRAM FOR
COMMUNICATION AMONG NODES IN A SYSTEM

BACKGROUND OF THE INVENTION

5 1. Field of the Invention

The present invention relates to a system, method, and program for enabling communication among nodes in a system.

2. Description of the Related Art

10 Machines or other devices may be comprised of a main controller node that communicates with specific hardware components using a bus interface, such as the Controller Area Network (CAN) serial bus. Each hardware component, e.g., a motor, etc., would include a CAN controller chip to allow communication with the CAN bus and other CAN devices on the bus. A main system controller can transmit commands and data to devices in the
15 system using the CAN message protocol. This arrangement has been satisfactory when there is no need for complex intercomponent communication, such as the case where a main system controller is the primary unit that manages and controls the electro-mechanical devices attached to the CAN bus.

As the complexity of processes that run within the devices and system controller
20 increase and the communication among such processes within the different devices increase, there is a need in the art for improved techniques for implementing the system components and for providing intercommunication among the processes executing within the system components.

25 SUMMARY OF THE PREFERRED EMBODIMENTS

Provided is a method, system, and program for allowing communication among nodes in a system. A request is received in a source node from a source object executing in the

source node to send a message to a destination object executing in a destination node. Each node includes a processor capable of multitasking multiple program objects and a communication interface to transmit and receive data with the other nodes. A determination is made in the source node as to whether the destination node and source node are a same node.

- 5 The message is sent in the source node to the destination object within the source node if the destination node is the source node. If the destination node is not the source node, then the source node transmits the message to the destination node through the communication interface. The destination node sends the message to the destination object within the destination node.

- In further embodiments, there is a message queue associated with each object in each
10 node. Sending the message to the destination object comprises invoking an operating system command to transmit the message to the message queue associated with the destination object.

- In still further embodiments, transmitting the message to the destination node comprises determining an address of the destination node that addresses the destination node when transmitting messages through the transmission medium. At least one message packet is
15 generated including the message, the destination node address, and an address of the destination object, and the at least one message packet is transmitted to the destination node.

- Still further, each node may be associated with one component of a system. In such case, a first node comprises a controller node, a second node comprises a component node that controls an electro-mechanical component of the system, the source object comprises a
20 work management object that manages system commands and the message includes a command to instruct a motion object in the component node to control the electro-mechanical component to perform an operation.

- Yet further, each object may be assigned a unique identifier in the system, wherein the unique identifier is used within all nodes to identify the destination object to receive the
25 message. Each node may also be assigned a unique node identifier used within all nodes to identify the destination node to receive the message.

In still further embodiments, a function call may receive the request from the source object to send the message to the destination object, determine whether the destination node is the same node, and then send the message to the destination object or cause the transmittal of the message to the destination node. The function call maintains the object and node identifier
5 assignment. In such cases, the node and object identifier used by each function call in each node may be updated to reflect a later modification to the arrangement of nodes or objects in the system.

The described implementations provide an improved system to enable program objects to communicate with other objects in a system comprised of multiple nodes. The system may
10 include nodes that control particular electro-mechanical components of the system. A source object would send a message to a destination object, and the routing of the message to a local queue for the destination object or to another node is seamlessly handled by a function call available at each node. In this way, the source object does not have to be concerned about the location of the destination object as the routing of the message to the destination object is
15 handled by a function call.

Moreover, in certain implementations, the function call uses a unique identifier for each node and object throughout the whole system when determining how to route a message to a destination object. In such embodiments, any modification to the arrangement of objects and nodes can be easily updated by merely updating the unique object and node address
20 assignment used by the function calls in each node.

BRIEF DESCRIPTION OF THE DRAWINGS

Referring now to the drawings in which like reference numbers represent corresponding parts throughout:

25 FIG. 1 is a block diagram illustrating of components in a tape library system known in the art;

FIG. 2 illustrates an architecture of processor nodes in a storage library system in accordance with preferred embodiments of the present invention;

FIG. 3 illustrates objects executing within the nodes in accordance with preferred embodiments of the present invention;

5 FIG. 4 illustrates fields in a message in accordance with preferred embodiments of the present invention;

FIG. 5 illustrates logic implemented in a send message function call to transmit messages between object in accordance with preferred embodiments of the present invention;

10 FIG. 6 illustrates logic to transmit a message across a transmission medium in accordance with preferred embodiments of the present invention; and

FIG. 7 illustrates logic to send a message received over the transmission medium to the destination object in accordance with preferred embodiments of the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

15 In the following description, reference is made to the accompanying drawings which form a part hereof and which illustrate several embodiments of the present invention. It is understood that other embodiments may be utilized and structural and operational changes may be made without departing from the scope of the present invention.

20 The Nodal Tape Library System

In certain cases, the node system is implemented in a storage library system that may include certain of the storage library components shown in FIG. 1. Such storage library components may include an array of storage cells, i.e., storage slots, that hold storage cartridges, such as optical or magnetic discs that are portable and removable from the library.

25 The term "storage cartridge" as used herein refers to any structure for housing such removable information media known in the art. With reference to FIG. 1, the storage library 2 includes a controller, an input/output station, a gripper assembly 4 capable of picking up and inserting

storage cartridges and an XY system 6 to move the gripper assembly 4 along the XY axis to a desired library element, such as storage slots 8a, b or drives 10a, b. Data operations are performed when the storage cartridge is inserted in the drives 10a, b. Once inserted in the drive 10a, b, data can be read from the cartridge by a host processor. Data transmitted from the host processor can be written to the storage cartridge inserted in the drive 10a, b. The library controller includes a processor, random access memory (RAM), and other controls and interfaces to direct the actions of the library components. The controller further interacts with a system to respond to library commands transmitted from the system. The input/output station is the opening through which the user may insert or remove a cartridge. An operator panel on the outside of the box 2 housing the tape library allows the user to communicate with the library controller. The library 2 also includes an access door 12 through which the user may add or remove cartridges maintained in the storage cells 8a, b.

The term "library element" as used herein refers to any location in the library 2 at which a storage cartridge may be located, e.g., the input/output stations, the storage cells 8a, b, the drives 10a, b, and gripper 4.

The gripper assembly 4 may also be equipped with a machine vision system, such as a bar code reader, to read a label on the storage cartridge when the gripper assembly 4 is positioned near another library element.

FIG. 2 illustrates an implementation of a nodal system within a storage library 20. The storage library 20 includes an accessor 22, which includes an XY system 24 and gripper assembly 26, an operator panel 28, and a host interface 30. The XY system 24 includes servo-electronics to move the gripper assembly 26 in the horizontal "X" and vertical "Y" directions to position the gripper assembly 26, which includes a robotic hand or picker as shown as element 4 in FIG. 1, to an appropriate storage slot or drive, such as the storage slots 8a, b and drives 10a, b shown in FIG. 1. The gripper assembly 26 may remove or insert a storage cartridge from a storage slot or drive. The gripper assembly 26 further includes a bar code scanner 32 which can read identification labels on the storage cartridges. The operator

panel 28 includes a display 34 to provide information to an operator and user interface controls to receive operator commands.

The host interface 30 provides an interface to host systems 36a, b, c over a communication line, such as a network, serial interface, Small Computer System Interface (SCSI), etc. Additionally, the host interface 30 may communicate with a separate web server or include an embedded web server to enable communication over a network, e.g., an Intranet, Local Area Network (LAN), the Internet, etc. The host systems 36a, b, c can communicate commands and receive data from the library 20 through the host interface 30. In further embodiments, the host systems 36a, b, c may communicate with the library 20 through data storage drives (not shown).

In preferred embodiments, each of the storage library components 24, 26, 28, and 30 include a processor node, an XY processor node 38, accessor processor node 40, operator panel processor node 42, and host communication processor node 44, respectively. Each processor node 38, 40, 42, and 44 comprises a processor, memory, firmware to control the processor, and a port to allow communication with a bus interface 46 through which the processor nodes communicate. The bus interface 46 may comprise a controller area network (CAN) bus known in the art, which is a multi-drop network, having a standard access protocol and wiring standards. In alternative embodiments, the bus 46 may comprise any bus or communication interface known in the art. Each of the processor nodes 38, 40, 42, and 44 may either recognize a message identifier associated with each message on the bus interface 46, in accordance with the CAN protocol, or may be specifically addressed with each message, for example as is known in the SCSI bus standard.

In CAN embodiments, each node 38, 40, 42, and 44 would include a CAN controller chip. Each node 38, 40, 42, and 44 would be assigned a unique CAN identifier and groups of nodes may also be assigned a unique identifier to allow messages to be broadcast to multiple nodes. Such CAN identifier may comprise a network address, unique number, etc. The CAN chip in each node 38, 40, 42, and 44 would be configured to receive messages including an

identifier or address assigned to address the particular node 38, 40, 42, and 44 and transmit messages including an identifier or address assigned to one of the other nodes 38, 40, 42, and 44. Details of programming the CAN chip are described in the publication entitled "82527 Serial Communications Controller Architectural Overview (Automotive)", having Intel order no. 5 272410-003 (Intel Corporation, Jan. 1996), which publication is incorporated herein by reference in its entirety.

In the described embodiments, the components of the library system 20 function as distributed computing elements, each operating under the control of their respective processor node, which performs system specific operations. In preferred embodiments, the accessor 10 processor node 40 functions as the central processor to receive, queue, execute, or distribute host system 36a, b, c commands. Thus, the accessor processor node 40 provides central processing facilities, including workflow management and queuing. In preferred embodiments, the host communication processor node 44 receives the host system 36a, b, c commands through the port and interface electronics provided by the host interface 30 and transfers the 15 commands to the accessor processor node 40 over the bus 46. The accessor processor node 40 can then execute such commands to control the gripper assembly 26 servo electronics to move the gripper assembly, or transfer XY motion commands to the XY processor node 38. The XY processor node 38 executes the commands to control the XY system 24 servo-electronics to move the accessor 22 in an XY direction through the tape library 20 to access a 20 storage cartridge in a drive or storage slot. The host system 36a, b, c command can instruct the accessor 22 to read the identification label on the storage cartridge at a particular location, access a storage cartridge at one library element (e.g., drive, storage slot, gripper, I/O slot, etc.) and move the storage cartridge to another library element.

25 In further embodiments, the library system 20 may include redundant instances of the above components 22, 28, and 30 to improve the availability of the system and increase processing capabilities. The co-pending and commonly assigned patent applications "High

Availability Work Queuing in an Automated Data Storage Library", having U.S. Application Serial No. 09/573,530, filed May 19, 2000 and "Automated Data Storage Library Distributed Control System", having U.S. Application Serial No. 09/573,531, filed May 19, 2000, which are both incorporated herein by reference in their entirety, describe further details of a tape
5 library system including distributed processor nodes to implement the components of the tape library system.

Nodal Communication

FIG. 3 illustrates further detail of software components that execute within the XY
10 processor node 38, accessor processor node 40, host communication processor node 44, and operator panel processor node 42. Each node 38, 40, 42, and 44 includes a real time operating system (RTOS) 50a, b, c, d, such as the RTOS described in the publication "ThreadX the high-performance embedded kernel: USER GUIDE" (Express Logic, Inc. 1997-1998), which publication is incorporated herein by reference in its entirety. Each node further
15 includes a CAN object 52a, b, c, d which manages message traffic between the node and the bus interface 46. The CAN objects 52a, b, c, d are capable of configuring the CAN chips included within the node.

Each node further includes objects which perform operations unique to that node. For instance, the accessor processor node 40 includes a work queue object 56 that queues and
20 executes tape library commands from the host systems 36a, b, c. The host communication processor node 44 includes a host communication object 58 that manages host system 36a, b, c (FIG. 2) communication and transmits host tape library commands to the work queue object 56 to queue and execute. The XY processor node 38 includes a motion command object 60 that receives motion related commands from the work queue object 56 and executes such
25 commands to control motion servo-electronics to move the accessor in the X-Y directions. Similarly, the accessor processor node 40 may include objects to manipulate the servo-electronics of the gripper assembly 26 (FIG. 2). The operator panel processor node 42

includes a display object 62 that processes display commands from the work queue object 56 and, in response, generates commands to render output on the display 34 (FIG. 2). In this way, the nodes 38, 40, 42, and 44 include a RTOS 50a, b, c, d, which allows code objects to execute and multitask to perform the particular operations of that node.

5 In preferred embodiments, an object executing in one of the nodes 38, 40, 42, and 44 would call a send message function 64a, b, c, and d to handle the transmission of a message to a local object or an object at a remote node. Thus, the same function call 64a, b, c, d is used to transmit a message to any object, whether it is on local or remote node. In this way, message communication for an object is seamless. In preferred embodiments, each object in
10 each node is assigned a unique identifier, referred to as an object identifier. The send message function 64a, b, c, d includes a global mapping of each object to an object identifier, each node to a node identifier, and each object identifier to a node identifier indicating the node in which the object executes. The unique identifier may comprise a unique address, number, code, etc. Thus, for every object identifier there is an associated node number indicating the node in which
15 the object identified by the object number executes. The same global mapping is used by each send message function 64a, b, c, d in every node 38, 40, 42, 44.

FIG. 4 illustrates the format of the fields in a message header 70 that is filled-in by the send message function 64a, b, c, d to transmit a message to another object. The message header 70 includes a destination node 72, destination object 74, source node 76, and source
20 object 78 fields indicating the source and destination objects and nodes for a message.

In preferred embodiments, the CAN objects 52a, b, c, d are capable of transmitting messages to nodes on the interface 46 using a standard network transmission protocol, such as Transmission Control Protocol/Internet Protocol (TCP/IP), Ethernet protocols, proprietary communication protocols, etc. to communicate over a CAN interface 46. The CAN objects
25 52a, b, c, d would include the capability to determine the CAN message addresses that correspond to the destination node 72 specified in the message header 70. The CAN message address is used to address the nodes on the CAN bus 46. The source object can specify a

group of destination objects. In such case, the destination node field 72 can specify all or a group of nodes that include the destination objects. The CAN object would then select a CAN message address that addresses all the destination nodes specified in field 72 of the message 70. The send message function 64a, b, c, d fills in the fields of the message header 70.

5 FIG. 5 illustrates the logic implemented in the send message function 64a, b, c, d to transmit messages from one object to another. At block 100, the send message function 64a, b, c, d is called by one of the objects 56, 58, 60, 62 with a message to send to one other object with a parameter including the destination object identifier 74. In response, the send message function 64a, b, c, d determines (at block 102) the source object that invoked the call
10 and the source node from which the call was made. The send message function 64a, b, c, d may obtain such source node and object identifier information from the message transmission or use operating system commands to determine the source object and node from which the message was invoked. The send message function 64a, b, c, d then determines from the global mapping (at block 104) the destination node that includes the destination object identifier
15 provided in the parameter of the call to the send message function 64a, b, c, d. The send message function 64a, b, c, d then inserts (at block 106) the determined identifier information in the fields 72, 74, 76, and 78 of the message header 70. If (at block 108) the destination node 72 is the same as the source node 76, then the send message function 64a, b, c, d calls (at block 110) the RTOS queue send function to queue the message in the message queue of the
20 destination object. In the ThreadX RTOS, the function to send a message to a message queue is the tx_queue_send function. In preferred embodiments, there is one message queue assigned to each object. In further embodiments, multiple objects may share one or more message queues, or multiple message queues may be assigned to one object. Otherwise, if (at block 108) the destination node is remote, then the send message function 64a, b, c, d calls (at block
25 112) the RTOS queue send function to queue the message with header 70 in the CAN object 52a, b, c, d message queue.

FIG. 6 illustrates logic implemented in the CAN objects 52a, b, c, d to process messages queued in their message queue. Control begins at block 120 with the CAN object 52a, b, c, d access a message in the message queue. In preferred embodiments, the CAN objects 52a, b, c, d would maintain a mapping of node identifiers to CAN addresses that address the nodes on the bus interface 46. At block 122, the CAN object 52a, b, c, d determines the CAN address of the destination node(s) indicated in field 72. The CAN object 52a, b, c, d then generates (at block 124) message packets including the message header 70, a sequence number of the packets, the total number of packets, and the determined CAN address of the destination node. The CAN chips on the nodes 38, 40, 42, 44 determine whether they should read a message transmitted on the bus interface 46. In preferred embodiments, the CAN objects 52a, b, c, d divide a message into multiple packets for transmittal. In alternative embodiments, the CAN object 52a, b, c, d may generate a single message packet. At block 126, the CAN object 52a, b, c, d generates a last packet in the sequence including a checksum on the contents of all the previous packets for error checking purposes. The message packets are then transmitted (at block 128) to the bus 46.

FIG. 7 illustrates logic implemented in the CAN objects 52a, b, c, d to handle the receipt of message transmitted over the bus 46. At block 150, the CAN object 52a, b, c, d accesses a message in one of the receive message buffers maintained in the CAN chip for the node. As discussed, the CAN chip for one of the nodes 38, 40, 44 would compare the CAN address in a message transmitted on the bus 46 with the addresses associated with receive message buffers in the CAN chip to determine whether the message on the bus 46 is addressed to that node 38, 40, 42, 44. If the message is addressed to that node 38, 40, 42, 44, then the CAN chip stores the message in the receive message buffer. Upon receipt of the last packet or after a timeout period, the CAN object 52a, b, c, d would then determine (at block 152) if all the packets in a sequence of packets for a message were received in the CAN chip. If not, then the CAN object 52a, b, c, d enters (at block 154) an error mode to recover the missing packets. Otherwise, if all packets were received, the CAN object 52a, b, c, d performs (at

block 156) an error checking operation on the packet using the checksum in the last packet to determine if the message data has been corrupted during transmission on the bus 46. If (at block 158) the packets do not pass the error checking test, then the CAN object 52a, b, c, d enters (at block 154) the error checking mode to request that the message be resent or some
5 other error recovery action. If the packets pass error checking, then the CAN object 52a, b, c, d calls (at block 160) the RTOS queue message function to queue the message in the destination object message queue indicated in the destination object 74 field of the message header 70 included in the message.

In further embodiments, the CAN objects 52a, b, c, d on each node 38, 40, 42, and
10 44 can transmit messages or pings to determine the availability of the other nodes and determine the immediate status on the connection to the destination object. In this way, the send message function 64a, b, c, d can be assured that the destination node is available when transmitting messages through the CAN objects 52a, b, c, d.

Preferred embodiments provide a technique for allowing objects in a distributed
15 computing environment in a system to communicate seamlessly using a single send message function. This single send message function determines whether to route the message to a local object message queue using a standard RTOS queue send function. If the object is not local, then the preferred embodiment send message function sends the message to an object to handle the transmission of the message over the bus interface.

20 With the preferred embodiments, if the arrangement of objects is altered, then all that has to be updated is the global mapping used by each send message function call. The updated global mapping indicates a new assignment of object identifiers to objects and association of object identifiers to nodes. For instance, if objects from one or more nodes are later consolidated into fewer nodes, then only the global mapping indicating the assignment of objects
25 to nodes would have to be updated in each node to allow the send message function call to use the new arrangement. Further, the removal or addition of nodes and objects can readily be indicated by updating the global mapping to reflect the new arrangement of nodes and objects.

Further, in preferred embodiments, each of the nodes functions independently of the other. This arrangement is different than many CAN systems where there is a master controller and various slaves devices connected to the CAN bus. With the preferred embodiments, each node is an independent processing unit and the preferred embodiments provide a

- 5 communication interface to allow processes and objects executing independently on the nodes to seamlessly communicate with objects executing in the same or different nodes in the systems.

- Following are some alternative implementations for the preferred embodiments. The preferred embodiments may be implemented as a method, apparatus or article of manufacture using standard programming and/or engineering techniques to produce software, firmware,
- 10 hardware, or any combination thereof. The term "article of manufacture" as used herein refers to code or logic implemented in hardware logic (e.g., an integrated circuit chip, Field Programmable Gate Array (FPGA), Application Specific Integrated Circuit (ASIC), etc.) or a computer readable medium (e.g., magnetic storage medium (e.g., hard disk drives, floppy disks,, tape, etc.), optical storage (CD-ROMs, optical disks, etc.), volatile and non-volatile
- 15 memory devices (e.g., EEPROMs, ROMs, PROMs, RAMs, DRAMs, SRAMs, firmware, programmable logic, etc.). Code in the computer readable medium is accessed and executed by a processor. The code in which preferred embodiments are implemented may further be accessible through a transmission media or from a file server over a network. In such cases, the article of manufacture in which the code is implemented may comprise a transmission media,
- 20 such as a network transmission line, wireless transmission media, signals propagating through space, radio waves, infrared signals, etc. Of course, those skilled in the art will recognize that many modifications may be made to this configuration without departing from the scope of the present invention, and that the article of manufacture may comprise any information bearing medium known in the art.

- 25 In preferred embodiments, the processor nodes comprised processors that operated under firmware control. In alternative embodiments, the processor nodes may be implemented as hardware logic, such as logic on an integrated circuit chip, e.g., an ASIC, FPGA, etc.

In preferred embodiments, a CAN object is used to generate a message compatible with the CAN protocol to transmit over a CAN bus, using a CAN messaging protocol and CAN chip. In this way, the CAN object provides the interface between the node and the CAN chip. In alternative embodiments, the bus interface may comprise any bus interface
5 known in the art, e.g., Ethernet, LAN etc., and the CAN object may comprise a network object to transmit the message using a network transmission protocol, e.g., TCP/IP, user datagram protocol (UDP), etc., other than CAN. In still further embodiments, the nodes may communicate using any transmission medium known in the art, including a wireless transmission medium. In such alternative embodiments, each node would include hardware to enable
10 communication through the transmission medium according to a communication system and protocol known in the art.

The preferred logic of FIGs. 5-7 describe specific operations occurring in a particular order. In alternative embodiments, certain of the logic operations may be performed in a different order, modified or removed and still implement preferred embodiments of the present
15 invention. Moreover, steps may be added to the above described logic and still conform to the preferred embodiments.

In preferred embodiments, the nodes were used to implement the components of a tape library system. In alternative embodiments, the nodes may be part of any system or machine known in the art, where the components of the machine or system may be implemented as
20 distributed processing nodes that communicate over a bus interface.

The foregoing description of the preferred embodiments of the invention has been presented for the purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise form disclosed. Many modifications and variations are possible in light of the above teaching. It is intended that the scope of the invention be limited
25 not by this detailed description, but rather by the claims appended hereto. The above specification, examples and data provide a complete description of the manufacture and use of the composition of the invention. Since many embodiments of the invention can be made

without departing from the spirit and scope of the invention, the invention resides in the claims hereinafter appended.

the first step of the process is to determine the scope of the invention. This is done by identifying the features of the invention and the prior art. The next step is to determine the novelty of the invention. This is done by comparing the features of the invention to the prior art. If the invention is novel, then it is eligible for a patent. The final step is to draft the patent claims. This is done by identifying the features of the invention and the prior art, and then drafting claims that cover the invention.